character of an earlier book by Lodge[4]). Penrose attacks the self from the perspective of literary theory. As Lodge put it in his earlier novel, she believes that there is no 'finite, unique soul or essence that constitutes a person's identity; there is only a subject position in an infinite web of discourses'. In his *Consciousness Explained*[5], Daniel Dennett approvingly quotes this passage, as an exaggerated version of his own view of the Self as a Centre of Narrative Gravity. Reed's talk is a reaction to these proposals, a defense of the Cartesian self.

It is a strange thing to say about a novel, but *Thinks…* is an excellent introduction to the modern scientific study of consciousness. At one point, Reed expresses her dismay that there are no literary people invited to present at an important conference on the topic – she sees this as a serious omission. After reading this provocative, thoughtful, and entirely enjoyable novel, it is obvious that – here at least – she is exactly right.

Paul Bloom
Yale University, Dept of Psychology, PO Box 208205, New Haven, CT 06520-8205, USA.
e-mail: Paul.Bloom@Yale.edu

References

1 Searle, J. (1980) Minds, brains, and programs. *Behav. Brain Sci.* 3, 417–457
2 Nagel, T. (1974) What is it like to be a bat? *Philos. Rev.* 83, 435–450
3 Jackson, F. (1986) What Mary didn't know. *J. Philos.* LXXXIII, 5, 291–295
4 Lodge, D. (1988) *Nice Work*, Sicker & Warburg
5 Dennett, D.C. (1991) *Consciousness Explained*, Little, Brown & Co.

Uncovering analogy

The Analogical Mind: Perspectives from Cognitive Science
edited by Dedre Gentner, Keith J. Holyoak and Boicho N. Kokinov, MIT Press, 2001.
$39.95 (xii + 541 pages)
ISBN 0 262 07206 8

In the summer of 1998, Boicho Kokinov organized a workshop at the New Bulgarian University that brought together leading researchers working in analogy and related fields. The edited volume *The Analogical Mind* derives from this workshop. The book reflects points of interest, convergence and controversy within the analogy community.

As I read the diverse collection of chapters, one basic question plagued me: what is an analogy? The stock reply to this question is that an analogy is a comparison of two mental representations that contain constituent parts. The comparison, or structural alignment, of these two representations establishes mappings or correspondences between them that respect their emergent structural similarities. To use Gentner's classic example[1], the Rutherford atom and the solar system can be seen as analogous because comparing the representations of these two domains results in the nucleus mapping to the sun and the electrons mapping to the planets. This analogical mapping is structurally sound (each part of one representation maps to one, and only one, part of the other representation) and highlights existing parallel structures (e.g. the electrons revolve around the nucleus and the planets revolve around the sun). Once satisfying correspondences are established, inferences about one domain can be drawn from knowledge of the other domain.

This stock reply proves to be satisfactory only at a general level when one considers the impressive range of problem domains (e.g. metaphor, mental simulation, problem solving, decision making) and methods (e.g. experimental, computational, developmental, sociological, linguistic) covered by the volume's contributors. At times I had the uneasy feeling that analogy subsumes all of cognition, explaining everything (or perhaps nothing). The concluding chapter by Hofstadter explicitly adopts the 'everything is analogy' position and suggests that all of vision is analogy. Of course, this statement is true at a trivial level (different retinal projections are 'analogous' to stored representations, thus enabling recognition), but this account de-emphasizes crucial domain-specific knowledge, processes and hardware. Construing analogy as smart pattern-matching subsumes not only vision, but also other broadly defined research areas, such as categorization research, which, like analogy, is concerned with how humans appreciate the sameness or equivalence of objects that are not identical.

In contrast to Hofstadter, other contributors seek to put limits on the application of analogy. For example, Keane and Costello argue that noun–noun phrases (i.e. conceptual combinations like 'bullet sprinter') are not interpreted through structural alignment. Although Hofstadter is trivially correct, Keane and Costello might be trivially incorrect because the crux of their argument seems to rest on the astute observation that conceptual combination uses constraints not typically embodied in analogical models. In other words, domain- and task-specific considerations need to be taken into account. This debate over the proper place of analogy is not merely semantic. For instance, it is not clear if the theoretical interpretation of the developmental and cross-species studies described in this volume depends on accepting analogy as a single 'thing'.

Perhaps partly in response to these issues, other chapters chart out alternative programs for understanding analogy. For example, Ken Forbus argues that analogy research should focus on the development of large-scale software systems that tackle real-world problems. Forbus thinks that such systems will challenge our assumptions and will stress the integration of analogy with other systems, thus revealing how analogy fits within larger frameworks. This line of research keeps an eye open towards specifying the qualitative nature of the algorithms, inputs, outputs and internal representations necessary to achieve the goals of a larger system.

Another line of research advocated is to ground analogy in the limitations of our cognitive architecture. Two chapters examine how working-memory limitations constrain analogical mapping. This line of research might eventually lead to models that capture a wide range of performance data from humans, and to an understanding of the neurological underpinnings of certain analogical tasks. Like Ken Forbus's work, these programs also have an inherent focus on integration. For example, Hummel and Holyoak consider the interrelations between analogy, working memory, higher-level vision and discourse processing.

The best way to view analogy might be as the 'glue' that enables integration. Certainly, the diverse nature of this book's chapters suggests that the principles underlying analogical comparisons are widespread in cognition. Paradoxically,

 1364-6613/01/$ – see front matter PII: S1364-6613(00)01747-2

the ultimate success of analogy research might lie in its demise (and infiltration into other research areas). Given the ubiquity of its principles and past examples of analogy cropping up in unexpected places (e.g. in artificial-grammar rule learning[2]), cognitive scientists would be well served by reading *The Analogical Mind*.

Bradley C. Love
Dept of Psychology, University of Texas at Austin, Austin, TX 78712, USA.
e-mail: love@psy.utexas.edu

References

1 Gentner, D. (1983) Structure-mapping: a theoretical framework for analogy. *Cognit. Sci.* 7, 155–170
2 Brooks, L.R. and Vokey, J.R. (1991) Abstract analogies and abstracted grammars: comments on Reber (1989) and Mathews *et al.* (1989). *J. Exp. Psychol. Gen.* 120, 316–323

Film Review

Artificiality embodied

AI: Artificial Intelligence
Directed by Steven Spielberg, starring Haley Joel Osment, Jude Law.
(UK release from September 21, 2001.)

Movies have long dealt with the humanization of machines and the mechanization of humans. Fritz Lang's seminal *Metropolis* (1927) brought us the robot Maria, who assumes human form and leads the working classes into revolt against 'the thinkers'. In *Demon Seed* (1977), Julie Christie tries (and fails) to avoid becoming the surrogate mother of an intelligent computer that becomes driven by its desire to procreate. Now, in *AI: Artificial Intelligence*, Steven Spielberg investigates the concept of thinking robots at a higher level, but only glances at it. Haley Joel Osment is stunning as David, who, we are told, is an 11-year-old boy whose love is real, although he himself is not. Robot David is given to a couple whose real son is cryogenically frozen after being comatose, and once the unique code words are said to David by his would-be mother, David loves his mother without condition. But all goes awry when the real son miraculously awakens, comes home and rejects David, eventually forcing his adoptive mother to dump him in the woods, unwittingly into the care of a robot male prostitute (Jude Law with a wax finish). With Pinnochio in mind, this unlikely synthetic pair heads off to search for the Blue Fairy, who can make David, and hopefully his love for Mommy, real.

From a scientific perspective, the film raised, but did not really get to grips with, the notion of artificial consciousness . This operated at several levels in the film: there was Teddy, the apparently indestructible supertoy, who gradually revealed more and more unbelievable

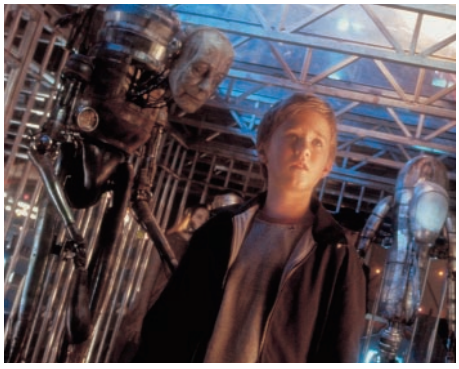

© 2001 Warner Bros/Dreamworks

abilities. An extra large dollop of anthropomorphism is clumsily thrown in here, as well as cognitive abilities so advanced in Teddy that David didn't seem to represent that much of an advance – just the addition of a 'love' subroutine. In many ways, *AI* calls to mind Searle's Chinese Room argument: can you ask questions of David that will allow you to distinguish whether he is robot or human? When he was rejected by Mommy, David's reactions were indistinguishable from those of a human child. Or, to pose the question in more emotive terms, how realistic does a robot have to be before you become unwilling to switch it off at night? In *Bladerunner* (1981), the question that sorted the men from the robots was what they would do if they found an upturned tortoise in the desert. *AI* takes a more subtle approach, but with ultimately the same outcome. The fact that Mommy could not bring herself to take David back to the factory to be dismantled indicates that she was, by this time, persuaded that he was more than a mechanical object. Sophistication is direly lacking in the moral issues raised in *AI* regarding artificial consciousness. There is the portrayal of human resistance to all the other robots in the film, in scenes in which a quasi-religious group aimed to 'rid the earth of artificial life'. Only workers in the robot business would be sympathetic to their plight. But the real questions here are, Who will be in charge of whom? How will we decide whether decisions taken by machines are morally good or morally bad?

Spielberg picked up this story from Stanley Kubrick, who had been developing it for some 18 years until his death in 1999. Unlikely bedfellows, Kubrick and Spielberg allegedly collaborated on this project for several years. But Kubrick's typically misanthropic vision is fully absent. AI's future is heavy-handedly polar – good people live in an Ikea world, wooden, chrome and neat: bad guys and robots occupy a kind of fiery neon heavy-metal gig. Despite Kubrick's influence, Spielberg can't resist giving in to his overwhelming fascination with schmaltz.

One of Kubrick's most enduring and popular characters was HAL, the supercomputer from *2001: A Space Odyssey* (1968), a film that dealt with the evolution of consciousness. HAL displayed genuine artificial intelligence, so much so, in fact, that his was the most rounded character in *2001*, complete with massive, murderous flaws. David is not as smart as HAL, and unfortunately his love is not real. And this is the fundamental problem with *AI*. It asks us to sympathize with a machine that has no understanding of why or how it has emotional cognitive functions in the first place. David's artificially intelligent traits are no different from those of Arnold Schwarzenegger's cyborg assassin, *Terminator*: he learns, and copies human behaviour and emotions. And that's not enough, neither for David, nor for an audience hoping to learn something about 'real' AI.

Adam Rutherford
adam.rutherford@current-trends.com
Julian Ogilvie
Assistant Editor, *TICS*

 1364-6613/01/$ – see front matter PII: S1364-6613(00)01775-7